arXiv:1611.08350v1 [cs.CV] 25 Nov 2016

Learning an Invariant Hilbert Space for Domain Adaptation

Samitha Herath SAMITHA.HERATH@DATA61.CSIRO.AU
College of Engineering and Computer Science
Australian National University & Data61, CSIRO
Canberra, Australia

Mehrtash Harandi MEHRTASH.HARANDI@DATA61.CSIRO.AU
College of Engineering and Computer Science
Australian National University & Data61, CSIRO
Canberra, Australia

Fatih Porikli FATIH.PORIKLI@DATA61.CSIRO.AU
College of Engineering and Computer Science
Australian National University & Data61, CSIRO
Canberra, Australia

Abstract

This paper introduces a learning scheme to construct a Hilbert space (i.e., a vector space along its inner product) to address both unsupervised and semi-supervised domain adaptation problems. This is achieved by learning projections from each domain to a latent space along the Mahalanobis metric of the latent space to simultaneously minimizing a notion of domain variance while maximizing a measure of discriminatory power. In particular, we make use of the Riemannian optimization techniques to match statistical properties (e.g., first and second order statistics) between samples projected into the latent space from different domains. Upon availability of class labels, we further deem samples sharing the same label to form more compact clusters while pulling away samples coming from different classes.We extensively evaluate and contrast our proposal against state-of-the-art methods for the task of visual domain adaptation using both handcrafted and deep-net features. Our experiments show that even with a simple nearest neighbor classifier, the proposed method can outperform several state-of-the-art methods benefitting from more involved classification schemes.

1. Introduction

This paper presents a learning algorithm to address both *unsupervised* (Gong et al., 2012; Fernando et al., 2013; Sun et al., 2016) and *semi-supervised* (Hoffman et al., 2014; Duan et al., 2012; Hubert Tsai et al., 2016) domain adaptation problems. Our goal here is to learn a latent space in which domain disparities are minimized. We show such a space can be learned by first matching the statistical properties of the projected domains (*e.g.*, covariance matrices), and then adapting the Mahalanobis metric of the latent space to the labeled data, *i.e.*, minimizing the distances between pairs sharing the same class label while pulling away samples with different class labels. We develop a geometrical solution to jointly learn projections onto the latent space and the Mahalanobis metric there by making use of the concepts of Riemannian geometry.

Thanks to deep learning, we are witnessing rapid growth in classification accuracy of the imaging techniques if substantial amount of labeled data is provided (Krizhevsky et al., 2012; Simonyan and Zisserman, 2014; He et al., 2016). However, harnessing the attained knowledge into a new application with limited labeled data (or even without having labels) is far beyond clear (Long et al., 2016; Ganin and Lempitsky, 2015; Tzeng et al., 2014; Chen et al., 2015). To make things

even more complicated, due to the inherit *bias* of datasets (Torralba and Efros, 2011; Shimodaira, 2000), straightforward use of large amount of auxiliary data does not necessarily assure improved performances. For example, the ImageNet (Russakovsky et al., 2015) data is hardly useful for an application designed to classify images captured by a mobile phone camera. Domain adaptation is the science of reducing such undesired effects in transferring knowledge from the available auxiliary resources to a new problem.

The most natural solution to the problem of DA is by identifying the structure of a common space that minimizes a notion of domain mismatch. Once such a space is obtained, one can design a classifier in it, hoping that the classifier will perform equally well across the domains as the domain mismatched is minimized. Towards this end, several studies assume that either **1.** a subspace of the target[1] domain is the right space to perform DA and learn how the source domain should be mapped onto it (Saenko et al., 2010; Hubert Tsai et al., 2016), or **2.** subspaces obtained from both source and target domains are equally important for classification, hence trying to either learn their evolution (Gopalan et al., 2011; Gong et al., 2012) or similarity measure (Shi et al., 2010; Wang and Mahadevan, 2011; Duan et al., 2012).

Objectively speaking, a common practice in many solutions including the aforementioned methods, is to simplify the learning problem by separating the two elements of it. That is, the algorithm starts by fixing a space (*e.g.*, source subspace of Fernando et al. (2013); Hubert Tsai et al. (2016)), and learns how to transfer the knowledge from domains accordingly. A curious mind may ask why should we resort to a predefined and fixed space in the first place.

In this paper, we propose a learning scheme that avoids such a separation. That is, we do not assume that a space or a transformation, apriori is known and fixed for DA. In essence, we propose to learn the structure of a Hilbert space (*i.e.*, its metric) along the transformations required to map the domains onto it jointly.

This is achieved through the following contributions,

(i) We propose to learn the structure of a latent space, along its associated mappings from the source and target domains to address both problems of unsupervised and semi-supervised DA.

(ii) Towards this end, we propose to maximize a notion of discriminatory power in the latent space. At the same time, we seek the latent space to minimize a notion of statistical mismatch between the source and target domains (see Fig. 1 for a conceptual diagram).

(iii) Given the complexity of the resulting problem, we provide a rigorous mathematical modeling of the problem. In particular, we make use of the Riemannian geometry and optimization techniques on matrix manifolds to solve our learning problem[2].

(iv) We extensively evaluate and contrast our solution against several baseline and state-of-the-art methods in addressing both unsupervised and semi-supervised DA problems.

1. In DA terminology target domain refers to the data directly related to the task. Source domain data is used as the auxiliary data for knowledge transferring.
2. Our implementation is available on `https://sherath@bitbucket.org/sherath/ils.git`.

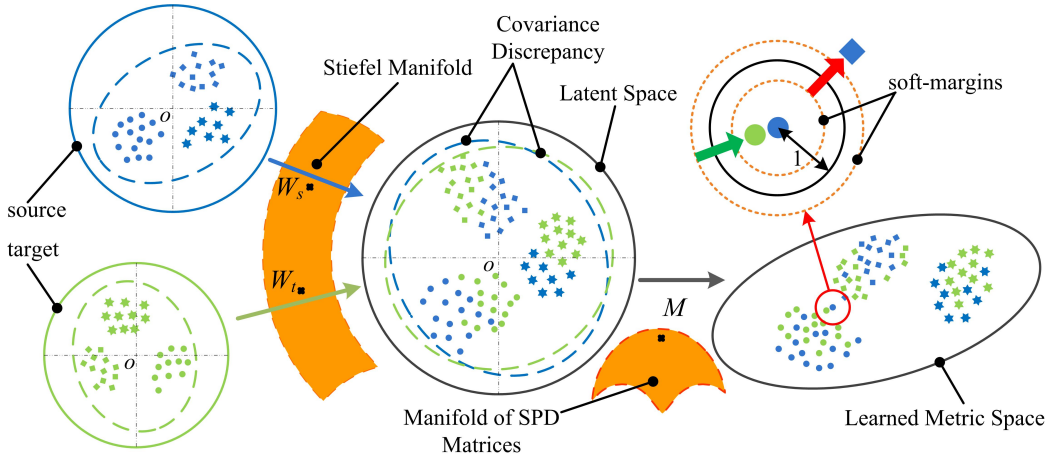

Figure 1: **A conceptual diagram of our proposal.** The marker shape represents the instance labels and color represents their original domains. Both source and target domains are mapped to a latent space using the transformations $\boldsymbol{W}_s$ and $\boldsymbol{W}_t$. The metric, $\boldsymbol{M}$ defined in the latent space is learned to maximize the discriminative power of samples in it. Indicated by dashed ellipsoids are the domain distributions. The statistical loss of our cost function aims to reduce such discrepancies within the latent space. Our learning scheme identifies the transformations $\boldsymbol{W}_s$ and $\boldsymbol{W}_t$ and the metric $\boldsymbol{M}$ jointly. This figure is best viewed in color.

2. Proposed Method

In this work, we are interested in learning an **I**nvariant **L**atent **S**pace (ILS) to reduce the discrepancy between domains. We first define our notations. Bold capital letters denote matrices (*e.g.*, $\boldsymbol{X}$) and bold lower-case letters denote column vectors (*e.g.*, $\boldsymbol{x}$). $\mathbf{I}_n$ is the $n \times n$ identity matrix. $\mathcal{S}_{++}^n$ and $\mathrm{St}(n,p)$ denote the SPD and Stiefel manifolds, respectively, and will be formally defined later. We show the source and target domains by $\mathcal{X}_s \subset \mathbb{R}^s$ and $\mathcal{X}_t \subset \mathbb{R}^t$. The training samples from the source and target domains are shown by $\{\boldsymbol{x}_i^s, y_i^s\}_{i=1}^{n_s}$ and $\{\boldsymbol{x}_i^t\}_{i=1}^{n_t}$, respectively. For now, we assume only source data is labeled. Later, we discuss how the proposed learning framework can benefit form the labeled target data.

Our idea in learning an ILS is to determine the transformations $\mathbb{R}^{s\times p} \ni \boldsymbol{W}_s : \mathcal{X}_s \to \mathcal{H}$ and $\mathbb{R}^{t\times p} \ni \boldsymbol{W}_t : \mathcal{X}_t \to \mathcal{H}$ from the source and target domains to a latent p-dimensional space $\mathcal{H} \subset \mathbb{R}^p$. We furthermore want to equip the latent space with a Mahalanobis metric, $\boldsymbol{M} \in \mathcal{S}_{++}^p$, to reduce the discrepancy between projected source and target samples (see Fig. 1 for a conceptual diagram).

To learn $\boldsymbol{W}_s$, $\boldsymbol{W}_t$ and $\boldsymbol{M}$ we propose to minimize a cost function in the form

$$\mathcal{L} = \mathcal{L}_d + \lambda \mathcal{L}_u \ . \tag{1}$$

In Eq. 1, $\mathcal{L}_d$ is a measure of dissimilarity between labeled samples. The term $\mathcal{L}_u$ quantifies a notion of statistical difference between the source and target samples in the latent space. As such, minimizing $\mathcal{L}$ leads to learning a latent space where not only the dissimilarity between labeled samples is reduced but also the domains are matched from a statistical point of view. The combination weight λ is envisaged to balance the two terms. The subscripts "d" and "u" in Eq. 1 stand for "Discriminative" and "Unsupervised". The reason behind such naming will become clear shortly. Below we detail out the form and properties of $\mathcal{L}_d$ and $\mathcal{L}_u$.

2.1 Discriminative Loss

The purpose of having $\mathcal{L}_d$ in Eq. 1 is to equip the latent space $\mathcal{H}$ with a metric to **1.** minimize dissimilarities between samples coming from the same class and **2.** maximizing the dissimilarities between samples from different classes.

Let $\boldsymbol{Z} = \{\boldsymbol{z}_j\}_{j=1}^{n}$ be the set of labeled samples in $\mathcal{H}$. In unsupervised domain adaptation $\boldsymbol{z}_j = \boldsymbol{W}_s^T \boldsymbol{x}_j^s$ and $n = n_s$. In the case of semi-supervised domain adaptation,

$$\boldsymbol{Z} = \left\{\boldsymbol{W}_s^T \boldsymbol{x}_j^s\right\}_{j=1}^{n_s} \bigcup \left\{\boldsymbol{W}_t^T \boldsymbol{x}_j^t\right\}_{j=1}^{n_{tl}},$$

where we assume n_{tl} labeled target samples are provided (out of available n_t samples). From the labeled samples in $\mathcal{H}$, we create N_p pairs in the form $(\boldsymbol{z}_{1,k}, \boldsymbol{z}_{2,k})$, $k = 1, 2, \cdots, N_p$ along their associated label $y_k \in \{-1, 1\}$. Here, $y_k = 1$ iff label of $\boldsymbol{z}_{1,k}$ is similar to that of $\boldsymbol{z}_{2,k}$ and -1 otherwise. That is the pair $(\boldsymbol{z}_{1,k}, \boldsymbol{z}_{2,k})$ is similar if $y_k = 1$ and dissimilar otherwise.

To learn the metric $\boldsymbol{M}$, we deem the distances between the similar pairs to be small while simultaneously making the distances between the dissimilar pairs large. In particular, we define $\mathcal{L}_d$ as,

$$\mathcal{L}_d = \frac{1}{N_p} \sum_{k=1}^{N_p} \ell_\beta\big(\boldsymbol{M}, y_k, \boldsymbol{z}_{1,k} - \boldsymbol{z}_{2,k}, 1\big) + r(\boldsymbol{M}), \tag{2}$$

with

$$\ell_\beta\big(\boldsymbol{M}, y, \boldsymbol{x}, u\big) = \frac{1}{\beta} \log\Big(1 + \exp\big(\beta y(\boldsymbol{x}^T \boldsymbol{M} \boldsymbol{x} - u)\big)\Big). \tag{3}$$

Here, ℓ_β is the generalized logistic function tailored with large margin structure (see Fig. 2) having a margin of u[3]. First note that the quadratic term in Eq. 3 (*i.e.*, $\boldsymbol{x}^T \boldsymbol{M} \boldsymbol{x}$) measures the Mahalanobis distance between $\boldsymbol{z}_{1,k}$ and $\boldsymbol{z}_{2,k}$ if used according to Eq. 2. Also note that $\ell_\beta\big(\cdot, \cdot, \boldsymbol{x}, \cdot\big) = \ell_\beta\big(\cdot, \cdot, -\boldsymbol{x}, \cdot\big)$, hence how samples are order in the pairs is not important.

To better understand the behavior of the function ℓ_β, assume the function is fed with a similar pair, *i.e.* $y_k = 1$. For the sake of discussion, also assume $\beta = 1$. In this case, ℓ_β is decreased if the distance between $\boldsymbol{z}_{1,k}$ and $\boldsymbol{z}_{2,k}$ is reduced. For a dissimilar pair (*i.e.*, $y_k = -1$), the opposite should happen to have a smaller objective. That is, the Mahalanobis distance between the samples of a pair should be increased.

The function $\ell_\beta\big(\cdot, \cdot, \boldsymbol{x}, \cdot\big)$ can be understood as a smooth and differentiable form of the hinge-loss function. In fact, $\ell_\beta\big(\cdot, \cdot, \boldsymbol{x}, \cdot\big)$ asymptotically reaches the hinge-loss function if $\beta \to \infty$. The smooth behavior of $\ell_\beta\big(\cdot, \cdot, \boldsymbol{x}, \cdot\big)$ is not only welcomed in the optimization scheme but also avoids samples in the latent space to collapse into a single point.

Along the general practice in metric learning, we regularize the metric $\boldsymbol{M}$ by $r(\boldsymbol{M})$. The divergences derived from the $\log\det(\cdot)$ function are familiar faces for regularizing Mahalanobis metrics in the litrature (Davis et al., 2007; Saenko et al., 2010).

Among possible choices, we make use of the Stein divergence (Cherian et al., 2013) in this work. Hence,

$$r(\boldsymbol{M}) = \frac{1}{p} \delta_s(\boldsymbol{M}, \boldsymbol{I_p}). \tag{4}$$

3. For now we keep the margin at $u = 1$ and later will use this to explain the soft-margin extension.

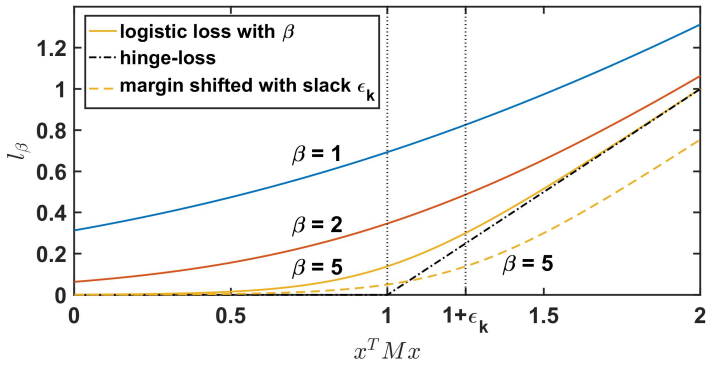

Figure 2: The behavior of the proposed ℓ_β(3) with $u = 1$ for various values of parameter β. Here, the horizontal axis is the value of the Mahalanobis distance and the function is plotted for $y = +1$. When $\beta \to \infty$, the function approaches the hinge-loss. An example of the soft-margin case (see Eq. 6), is also plotted for $\beta = 5$ curve. The figure is best seen in color.

Where,

$$\delta_s(\boldsymbol{P}, \boldsymbol{Q}) = \log\det\left(\frac{\boldsymbol{P}+\boldsymbol{Q}}{2}\right) - \frac{1}{2}\log\det\left(\boldsymbol{P}\boldsymbol{Q}\right), \tag{5}$$

for $\boldsymbol{P}, \boldsymbol{Q} \in \mathcal{S}_{++}$.

Our motivation in using the Stein divergence stems from its unique properties. Among others, Stein divergence is symmetric, invariant to affine transformation and closely related to geodesic distances on the SPD manifold (Cherian et al., 2013; Harandi et al., 2016).

Soft Margin Extension

For large values of β, the cost in Eq. 2 seeks the distances of similar pairs to be less than 1 while simultaneously it deems the distances of dissimilar pairs to exceed 1. This hard-margin in the design of $\ell_\beta\big(\cdot, \cdot, \boldsymbol{x}, \cdot\big)$ is not desirable. For example, with a large number of pairs, it is often the case to have outliers. Forcing outliers to fit into the hard margins can result in overfitting. As such, we propose a soft-margin extension of Eq. 3. The soft-margins are implemented by associating a non-negative slack variable ϵ_k to a pair according to

$$\mathcal{L}_d = \frac{1}{N_p}\sum_{k=1}^{N_p} \ell_\beta\big(\boldsymbol{M}, y_k, \boldsymbol{z}_{1,k} - \boldsymbol{z}_{2,k}, 1 + y_k\epsilon_k\big) + r(\boldsymbol{M}) + \frac{1}{N_p}\sqrt{\sum \epsilon_k^2}, \tag{6}$$

where a regularizer on the slack variables is also envisaged.

Matching Statistical Properties

A form of incompatibility between domains is due to their statistical discrepancies. Matching the first order statistics of two domains for the purpose of adaptation is studied in Pan et al. (2011); Baktashmotlagh et al. (2016); Hubert Tsai et al. (2016)[4]. In our framework, matching domain

4. The use of Maximum Mean Discrepancy (MMD) (Borgwardt et al., 2006) for domain adaptation is a well-practiced idea in the literature (see for example Pan et al. (2011); Baktashmotlagh et al. (2016); Hubert Tsai et al. (2016)). Em-

averages can be achieved readily. In particular, let $\bar{\boldsymbol{x}}_i^s = \boldsymbol{x}_i^s - \boldsymbol{m}_s$ and $\bar{\boldsymbol{x}}_j^t = \boldsymbol{x}_j^t - \boldsymbol{m}_t$ be the centered source and target samples with $\boldsymbol{m}_s$ and $\boldsymbol{m}_t$ being the mean of corresponding domains. It follows easily that the domain means in the latent space are zero[5] and hence matching is achieved.

To go beyond first order statistics, we propose to match the second order statistics (*i.e.*, covariance matrices) as well. The covariance of a domain reflects the relationships between its features. Hence, matching covariances of source and target domains in effect improves the cross feature relationships. We capture the mismatch between source and target covariances in the latent space using the $\mathcal{L}_u$ loss in Eq. 1. Given the fact that covariance matrices are points on the SPD manifold, we make use of the Stein divergence to measure their differences. This leads us to define $\mathcal{L}_u$ as

$$\mathcal{L}_u = \frac{1}{p}\delta_s(\boldsymbol{W}_s^T\boldsymbol{\Sigma}_s\boldsymbol{W}_s, \boldsymbol{W}_t^T\boldsymbol{\Sigma}_t\boldsymbol{W}_t), \tag{7}$$

with $\boldsymbol{\Sigma}_s \in \mathcal{S}_{++}^s$ and $\boldsymbol{\Sigma}_t \in \mathcal{S}_{++}^t$ being the covariance matrices of the source and target domains, respectively. We emphasize that matching the statistical properties as discussed above is an unsupervised technique, enabling us to address unsupervised DA.

3. Optimization

The objective of our algorithm is to learn the transformation parameters ($\boldsymbol{W}_s$ and $\boldsymbol{W}_t$), the metric $\boldsymbol{M}$ and slack variables $\epsilon_1, \epsilon_2, ...\epsilon_{N_p}$ (see Eq. 6 and Eq. 7). Inline with the general practice of dimensionality reduction, we propose to have orthogonality constraints on $\boldsymbol{W}_s$ and $\boldsymbol{W}_t$. That is $\boldsymbol{W}_s^T\boldsymbol{W}_s = \boldsymbol{W}_t^T\boldsymbol{W}_t = \mathbf{I}_p$. We elaborate how orthogonality constraint improves the discriminatory power of the proposed framework later in our experiments.

The problem depicted in Eq. 1 is indeed a non-convex and constrained optimization problem. One may resort to the method of Projected Gradient Descent (PGD) (Boyd and Vandenberghe, 2004) to minimize Eq. 1. In PGD, optimization is proceed by projecting the gradient-descent updates onto the set of constraints. For example, in our case, we can first update $\boldsymbol{W}_s$ by ignoring the orthogonality constraint on $\boldsymbol{W}_s$ and then project the result onto the set of orthogonal matrices using eigen-decomposition. As such, optimization can be performed by alternatingly updating $\boldsymbol{W}_s$, $\boldsymbol{W}_t$, the metric $\boldsymbol{M}$ and slack variables using PGD.

In PGD, to perform the projection, the set of constraints needs to be closed though in practice one can resort to open sets. For example, the set of SPD matrices is open though one can project a symmetric matrix onto this set using eigen-decomposition.

Empirically, PGD showed an erratic and numerically unstable behavior in addressing our problem. This can be attributed to the non-linear nature of Eq. 1, existence of open-set constraints in the problem or perhaps the combination of both. To alleviate the aforementioned difficulty, we propose a more principle approach to minimize Eq. 1 by making use of the Riemannian optimization technique. We take a short detour and briefly describe the Riemannian optimization methods below.

---

pirically, determining MMD boils down to computing the distance between domain averages when domain samples are lifted to a reproducing kernel Hilbert space. Some studies claim matching the first order statistics is a weaker form of domain adaptation through MMD. We do not support this claim and hence do not see our solution as a domain adaptation method by minimizing the MMD.

5. We note that $\sum \boldsymbol{W}_s^T\bar{\boldsymbol{x}}_i^s = \boldsymbol{W}_s^T\sum\bar{\boldsymbol{x}}_i^s = \mathbf{0}$. This holds for the target domain as well.

Optimization on Riemannian manifolds.

Consider a non-convex constrained problem in the form

$$\begin{aligned} &\text{minimize } f(\boldsymbol{x}) \\ &\text{s.t.} \quad \boldsymbol{x} \in \mathcal{M} \,, \end{aligned} \tag{8}$$

where $\mathcal{M}$ is a Riemannian manifold, *i.e.*, informally, a smooth surface that locally resembles a Euclidean space. Optimization techniques on Riemannian manifolds (*e.g.*, Conjugate Gradient) start with an initial solution $\boldsymbol{x}^{(0)} \in \mathcal{M}$, and iteratively improve the solution by following the geodesic identified by the gradient. For example, in the case of Riemannian Gradient Descent Method (RGDM), the updating rule reads

$$\boldsymbol{x}^{(t+1)} = \tau_{\boldsymbol{x}^{(t)}}\big(-\alpha \text{ grad } f(\boldsymbol{x}^{(t)})\big) \,, \tag{9}$$

with $\alpha > 0$ being the algorithm's step size. Here, $\tau_{\boldsymbol{x}}(\cdot) : T_{\boldsymbol{x}}\mathcal{M} \to \mathcal{M}$, is called the retraction[6] and moves the solution along the descent direction while assuring that the new solution is on the manifold $\mathcal{M}$, *i.e.*, it is within the constraint set. $T_{\boldsymbol{x}}\mathcal{M}$ is the tangent space of $\mathcal{M}$ at $\boldsymbol{x}$ and can be thought of as a vector space with its vectors being the gradients of all functions defined on $\mathcal{M}$.

We defer more details on Riemannian optimization techniques to the **Appendix A**. As for now, it suffices to say that to perform optimization on the Riemannian manifolds, the form of Riemannian gradient, retraction and the gradient of the objective with respect to its parameters (shown by ∇) are required. The constraints in Eq.1 are orthogonality (transformations $\boldsymbol{W}_s$ and $\boldsymbol{W}_t$) and p.d. for metric $\boldsymbol{M}$. The geometry of these constraints are captured by the Stiefel (James, 1976) and SPD (Harandi et al., 2016) manifolds, formally defined as

Definition 1 (The Stiefel Manifold) *The set of $(n \times p)$-dimensional matrices, $p \leq n$, with orthonormal columns endowed with the Frobenius inner product[7] forms a compact Riemannian manifold called the Stiefel manifold* $\text{St}(p, n)$ *(Absil et al., 2009).*

$$\text{St}(p, n) \triangleq \{\boldsymbol{W} \in \mathbb{R}^{n \times p} : \boldsymbol{W}^T\boldsymbol{W} = \mathbf{I}_p\} \,. \tag{10}$$

Definition 2 (The SPD Manifold) *The set of $(p \times p)$ dimensional real, SPD matrices endowed with the Affine Invariant Riemannian Metric (AIRM) (Pennec et al., 2006) forms the SPD manifold $\mathcal{S}^p_{++}$.*

$$\mathcal{S}^p_{++} \triangleq \{\boldsymbol{M} \in \mathbb{R}^{p \times p} \colon \boldsymbol{v}^T\boldsymbol{M}\boldsymbol{v} > 0, \forall \boldsymbol{v} \in \mathbb{R}^p - \{\mathbf{0}_p\}\}. \tag{11}$$

Updating $\boldsymbol{W}_s$, $\boldsymbol{W}_t$ and $\boldsymbol{M}$ and slacks can be done alternatively using Riemannian optimization. As mentioned above, the ingredients for doing so are **1.** the Riemannian tools for the Stiefel and SPD manifolds along **2.** the form of gradients of the objective with respect to its parameters. To do complete justice, in Table. 1 we provide the Riemmanian metric, form of Riemannian gradient and retraction for the Stiefel and SPD manifolds. In Table. 2, the gradient of Eq. 1 with respect to $\boldsymbol{W}_s$, $\boldsymbol{W}_t$ and $\boldsymbol{M}$ and slacks is provided. The detail of derivations can be found in the **Appendix**

6. Strictly speaking and in contrast with the exponential map, a retraction only guarantees to pull a tangent vector on the geodesic locally, *i.e.*, close to the origin of the tangent space. Retractions, however, are typically easier to compute than the exponential map and have proven effective in Riemannian optimization (Absil et al., 2009).
7. Note that the literature is divided between this choice and another form of Riemannian metric. See Edelman et al. (1998) for details.

Table 1: Riemannian metric, gradient and retraction on $\mathrm{St}(p,n)$ and $\mathcal{S}_{++}^{p}$. Here, $\mathrm{uf}(\boldsymbol{A}) = \boldsymbol{A}(\boldsymbol{A}^T\boldsymbol{A})^{-1/2}$, which yields an orthogonal matrix, $\mathrm{sym}(\boldsymbol{A}) = \frac{1}{2}(\boldsymbol{A}+\boldsymbol{A}^T)$ and $\mathrm{expm}(\cdot)$ denotes the matrix exponential.

	$\mathrm{St}(p,n)$	$\mathcal{S}_{++}^{p}$
Matrix representation	$\boldsymbol{W} \in \mathbb{R}^{n\times p}$	$\boldsymbol{M} \in \mathbb{R}^{p\times p}$
Riemannian metric	$g_\nu(\xi,\varsigma) = \mathrm{Tr}(\xi^T\varsigma)$	$g_{\mathcal{S}}(\xi,\varsigma) = \mathrm{Tr}\left(\boldsymbol{M}^{-1}\xi\boldsymbol{M}^{-1}\varsigma\right)$
Riemannian gradient	$\nabla_{\boldsymbol{W}}(f) - \boldsymbol{W}\mathrm{sym}\Big(\boldsymbol{W}^T\nabla_{\boldsymbol{W}}(f)\Big)$	$\boldsymbol{M}\mathrm{sym}\Big(\nabla_{\boldsymbol{M}}(f)\Big)\boldsymbol{M}$
Retraction	$\mathrm{uf}(\boldsymbol{W}+\xi)$	$\boldsymbol{M}^{\frac{1}{2}}\,\mathrm{expm}(\boldsymbol{M}^{-\frac{1}{2}}\xi\boldsymbol{M}^{-\frac{1}{2}})\boldsymbol{M}^{\frac{1}{2}}$

Table 2: Gradients of soft-margin ℓ_β and $\mathcal{L}_u$ w.r.t. the model parameters and slack variables. Without less of generality we only consider a labeled similar ($y_k = +1$) pair $\boldsymbol{x}_i^s$ and $\boldsymbol{x}_j^t$. Here, $r = \exp\left(\beta\big((\boldsymbol{W}_s^T\boldsymbol{x}_i^s - \boldsymbol{W}_t^T\boldsymbol{x}_i^t)^T\boldsymbol{M}(\boldsymbol{W}_s^T\boldsymbol{x}_i^s - \boldsymbol{W}_t^T\boldsymbol{x}_i^t) - 1 - e^{v_k}\big)\right)$.

$\nabla_{\boldsymbol{W}_s}\ell_\beta$	$\frac{2}{N_p}(1+r^{-1})^{-1}\boldsymbol{x}_i^s(\boldsymbol{x}_i^{sT}\boldsymbol{W}_s - \boldsymbol{x}_j^{t\,T}\boldsymbol{W}_t)\boldsymbol{M}$
$\nabla_{\boldsymbol{W}_t}\ell_\beta$	$\frac{2}{N_p}(1+r^{-1})^{-1}\boldsymbol{x}_j^t(\boldsymbol{x}_j^{t\,T}\boldsymbol{W}_t - \boldsymbol{x}_i^{sT}\boldsymbol{W}_s)\boldsymbol{M}$
$\nabla_{\boldsymbol{M}}\ell_\beta$	$\frac{1}{N_p}(1+r^{-1})^{-1}\big(\boldsymbol{W}_s^T\boldsymbol{x}_i^s - \boldsymbol{W}_t^T\boldsymbol{x}_j^t\big)\big(\boldsymbol{x}_i^{sT}\boldsymbol{W}_s - \boldsymbol{x}_j^{t\,T}\boldsymbol{W}_t\big)$
$\nabla_{v_k}\ell_\beta$	$\frac{-1}{N_p}e^{v_k}(1+r^{-1})^{-1}$
$\nabla_{\boldsymbol{W}_s}\mathcal{L}_u$	$\frac{1}{p}\boldsymbol{\Sigma}_s\boldsymbol{W}_s\Big(2\big(\boldsymbol{W}_s^T\boldsymbol{\Sigma}_s\boldsymbol{W}_s + \boldsymbol{W}_t^T\boldsymbol{\Sigma}_t\boldsymbol{W}_t\big)^{-1} - \big(\boldsymbol{W}_s^T\boldsymbol{\Sigma}_s\boldsymbol{W}_s\big)^{-1}\Big)$

B. A tiny note about the slacks worth mentioning. To preserve the non-negativity constraint on ϵ_k, we define $\epsilon_k = e^{v_k}$ and optimize on v_k instead. This in turn makes optimization for the slacks an unconstrained problem.

Remark 3 *From a geometrical point of view, we can make use of the product topology of the parameter space to avoid alternative optimization. More specifically, the set*

$$\mathcal{M}_{prod.} = \mathrm{St}(p,s) \times \mathrm{St}(p,t) \times \mathcal{S}_{++}^{p} \times \mathbb{R}^{N_p}, \tag{12}$$

can be given the structure of a Riemannian manifold using the concept of product topology (Absil et al., 2009).

Remark 4 *In Fig. 3, we compare the convergence behavior of PGD, alternating Riemannian optimization and optimization using the product geometry. While optimization on $\mathcal{M}_{prod.}$ convergences faster, the alternating method results in a lower loss. This behavior resembles the difference between the stochastic gradient descent compared to its batch counterpart.*

Remark 5 *The complexity of the optimization depends on the number of labeled pairs. One can always resort to a stochastic solution by sampling from the set of similar/dissimilar pairs if addressing a very large-scale problem. In our experiments, we did not face any difficulty optimizing with an i7 desktop machine with 32GB of memory.*

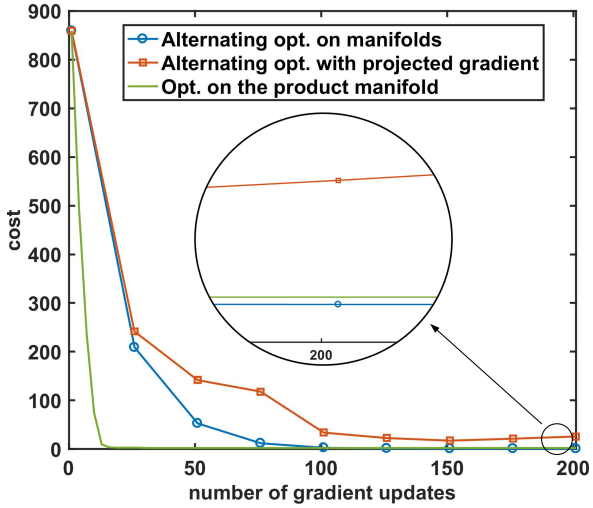

Figure 3: Optimizing Eq. 1 using PGD (red curve), Riemannian gradient descent using alternating approach (blue curve) and product topology (green curve). Optimization using the product topology converges faster but a lower cost can be attained using alternating Riemannian optimization.

4. Related Work

The literature on domain adaptation spans a very broad range (see Patel et al. (2015) for a recent survey). Our solution falls under the category of domain adaptation by subspace learning (DA-SL). As such, we confine our review only to methods under the umbrella of DA-SL.

One notable example of constructing a latent space is the work of Daumé III et al. (2010). In particular, the authors propose to use two fixed and predefined transformations to project source and target data to a common and higher-dimensional space. As a requirement, the method only accepts domains with the same dimensionality and hence cannot be directly used to adapt heterogeneous domains.

Gopalan et al. (2011) observed that the geodesic connecting the source and target subspaces conveys useful information for DA and proposed the Sampling Geodesic Flow (SGF) method (Gopalan et al., 2011). The Geodesic Flow Kernel (GFK) is an improvement over the SGF technique where instead of sampling a few points on the geodesic, the whole curve is used for domain adaptation (Gong et al., 2012). In both methods, the domain subspaces are fixed and obtained by Principal Component Analysis (PCA) or Partial Least Square regression (PLS) (Krishnan et al., 2011). In contrast to our solution, in SGF and GFK learning the domain subspaces is disjoint from the knowledge transfer algorithm. In our experiments, we will see that the subspaces determined by our method can even boost the performance of GFK, showing the importance of joint learning of domain subspaces and the knowledge transfer scheme.

Domain adaptation by fixing the subspace/representation of one of the domains is a popular theme in many recent works, as it simplifies the learning scheme. Examples are the max-margin adaptation (Hoffman et al., 2014; Duan et al., 2012), the metric/similarity learning of Saenko et al. (2010) and its kernel extension (Kulis et al., 2011), the landmark approach of Hubert Tsai et al. (2016), the alignment technique of Fernando et al. (2013), correlation matching of Sun et al. (2016)

and methods that use maximum mean discrepancy (MMD) (Borgwardt et al., 2006) for domain adaptation (Pan et al., 2011; Baktashmotlagh et al., 2016).

In contrast to the above methods, some studies opt to learn the domain representation along the knowledge transfer method jointly. Two representative works are the HeMap (Shi et al., 2010) and manifold alignment (Wang and Mahadevan, 2011). The HeMap learns two projections to minimize the instance discrepancies (Shi et al., 2010). The problem is however formulated such that equal number of source and target instances is required to perform the training. The manifold alignment algorithm of Wang and Mahadevan (2011) attempts to preserve the label structure in the latent space. However, it is essential for the algorithm to have access to labeled data in both source and target domains.

Our solution learns all transformations to the latent space. We do not resort at subspace representations learned disjointly to the DA framework. With this use of the latent space, our algorithm is not limited for applications where source and target data have similar dimensions or structure.

5. Experimental Evaluations

We run extensive experiments on both semi-supervised and unsupervised settings, spanning from the handcrafted features (SURF) to the current state-of-the-art deep-net features (VGG-Net). For comparisons, we use the implementations made available by the original authors. Our method is denoted by **ILS**.

5.1 Implementation Details

Since the number of dissimilar pairs is naturally larger than the number of similar pairs, we randomly sample from the different pairs to keep the sizes of these two sets equal. We initialize the projection matrices $\boldsymbol{W_s}$, $\boldsymbol{W_t}$ with PCA, following the transductive protocol Gong et al. (2012); Fernando et al. (2013); Hoffman et al. (2014); Hubert Tsai et al. (2016). For the semi-supervised setting, we initialize $\boldsymbol{M}$ with the Mahalanobis metric learned on the similar pair covariances (Köstinger et al., 2012), and for the unsupervised setting, we initialize it with the identity matrix. For all our experiments we have $\lambda = 1$. We include an experiment showing our solution's robustness to λ in the supplementary material. We use the toolbox provided by Boumal et al. (2014) for our implementations.

Remark 6 *To have a simple way of determining β in Eq. 3, we propose a heuristic which is shown to be effective in our experiments. To this end, we propose to set β to the reciprocal of the standard deviation of the similar pair distances.*

5.2 Semi-supervised Setting

In our semi-supervised experiments, we follow the standard setup on the Office+Caltech10 dataset with the train/test splits provided by Hoffman et al. (2013). The Office+Caltech10 dataset contains images collected from 4 different sources (see Fig. 4) and 10 object classes. The corresponding domains are **A**mazon, **W**ebcam, **D**SLR, and **C**altech. We use a subspace of dimension 20 for DASL algorithms. We employ SURF (Bay et al., 2006) for the handcrafted feature experiments. We

Figure 4: Samples from the Office+Caltech10 dataset (Saenko et al., 2010; Gong et al., 2012). Although the DSLR and Webcam images depict similar content, they have considerably different resolutions.

extract VGG-Net features with the network model of (Simonyan and Zisserman, 2014) for the deep-net feature experiments[8]. We compare our performance with the following benchmarks:

1-NN-t and SVM-t : Basic Nearest Neighbor (1-NN) and linear SVM classifiers trained only on the target domain.

HFA (Duan et al., 2012) : This method employs latent space learning based on the max-margin framework. As in its original implementation, we use the RBF kernel SVM for its evaluation.

MMDT (Hoffman et al., 2014) : This method jointly learns a transformation between the source and target domains along a linear SVM for classification.

CDLS (Hubert Tsai et al., 2016) : This is the cross-domain landmark search algorithm. We use the parameter setting ($\delta = 0.5$ in the notation of Hubert Tsai et al. (2016)) recommended by the authors.

Table 3, Table 4 and Table 5 report the performances using the handcrafted SURF, VGG-FC6 and VGG-FC7 features, respectively. For the SURF features our solution achieves the best performance in 7 out 12 cases. For the VGG-FC6 and VGG-FC7 features, our solution tops in 10 and 7 sets respectively. It seems that in comparison to VGG-FC6, VGG-FC7 features are less discriminative for all the DA algorithms. We notice the **1-NN-t** baseline performs the worst for both SURF and the VGG-Net features. Hence, it is clear that the used features do not favor the nearest neighbor classifier. We observe that **C**altech and **A**mazon domains contain the largest number of test instances. Although the performances of all tested methods decrease on these domains, particularly on **C**altech, our method achieves the top rank in almost all domain transformations.

8. The same SURF and VGG-Net features are used for the unsupervised experiments as well.

Table 3: Semi-supervised domain adaptation results using SURF features on Office+Caltech10 (Gong et al., 2012) dataset with the evaluation setup of (Hoffman et al., 2014). The best score (in bold blue), the second best (in blue).

	A→W	A→D	A→C	W→A	W→D	W→C	D→A	D→W	D→C	C→A	C→W	C→D	**Avg.**
1-NN-t	34.5	33.6	19.7	29.5	35.9	18.9	27.1	33.4	18.6	29.2	33.5	34.1	29.0
SVM-t	63.7	57.2	32.2	46.0	56.5	29.7	45.3	62.1	32.0	45.1	60.2	56.3	48.9
HFA	57.4	55.1	31.0	**56.5**	56.5	29.0	42.9	60.5	30.9	43.8	58.1	55.6	48.1
MMDT	64.6	56.7	36.4	47.7	67.0	32.2	46.9	74.1	34.1	49.4	63.8	56.5	52.5
CDLS	**68.7**	**60.4**	35.3	51.8	60.7	33.5	50.7	68.5	34.9	50.9	**66.3**	**59.8**	53.5
ILS (1-NN)	59.7	49.8	**43.6**	54.3	**70.8**	**38.6**	**55.0**	**80.1**	**41.0**	**55.1**	62.9	56.2	**55.6**

Table 4: Semi-supervised domain adaptation results using VGG-FC6 features on Office+Caltech10 (Gong et al., 2012) dataset with the evaluation setup of Hoffman et al. (2014). The best (in bold blue), the second best (in blue).

	A→W	A→D	A→C	W→A	W→D	W→C	D→A	D→W	D→C	C→A	C→W	C→D	**Avg.**
1-NN-t	81.0	79.1	67.8	76.1	77.9	65.2	77.1	81.7	65.6	78.3	80.2	77.7	75.6
SVM-t	89.1	**88.2**	77.3	86.5	87.7	76.3	87.3	88.3	76.3	87.5	87.8	84.9	84.8
HFA	87.9	87.1	75.5	85.1	87.3	74.4	85.9	86.9	74.8	86.2	86.0	**87.0**	83.7
MMDT	82.5	77.1	78.7	84.7	85.1	73.6	83.6	86.1	71.8	85.9	82.8	77.9	80.8
CDLS	**91.2**	86.9	78.1	87.4	88.5	78.2	88.1	90.7	77.9	88.0	89.7	86.3	85.9
ILS (1-NN)	90.7	87.7	**83.3**	**88.8**	**94.5**	**82.8**	**88.7**	**95.5**	**81.4**	**89.7**	**91.4**	86.9	**88.5**

Table 5: Semi-supervised domain adaptation results using VGG-FC7 features on Office+Caltech10 (Gong et al., 2012) dataset with the evaluation setup of Hoffman et al. (2014). The best (in bold blue), the second best (in blue).

	A→W	A→D	A→C	W→A	W→D	W→C	D→A	D→W	D→C	C→A	C→W	C→D	**Avg.**
1-NN-t	81.8	78.2	68.3	77.8	77.6	67.4	78.1	81.5	66.9	79.0	80.6	77.4	76.2
SVM-t	87.5	**85.4**	76.8	86.2	85.6	75.8	87.0	87.1	76.0	87.1	86.4	84.4	83.8
HFA	86.6	85.3	75.2	84.9	85.5	74.8	85.8	86.5	75.1	86.0	85.3	**84.8**	83.0
MMDT	76.9	73.3	78.1	83.6	79.5	72.2	82.3	83.8	71.7	85.3	77.8	72.6	78.1
CDLS	**90.0**	85.0	78.5	87.2	86.5	79.0	**87.7**	89.5	**78.8**	87.8	**89.7**	84.6	85.4
ILS (1-NN)	89.3	84.0	**81.9**	**88.4**	**91.0**	**80.8**	86.9	**94.4**	**78.8**	**88.9**	88.7	83.3	**86.4**

5.3 Unsupervised Setting

In the unsupervised domain adaptation problem, only labeled data from the source domain is available (Fernando et al., 2013; Gong et al., 2012). We perform two sets of experiments for this setting. (1) We evaluate the object recognition performance on the Office+Caltech10 dataset. Similar to the semi-supervised settings, we use the SURF and VGG-Net features. Our results demonstrate that the learned transformations by our method are superior domain representations. (2) We analyze our performance when the domain discrepancy is gradually increased. This experiment is performed on the PIE-Face dataset. We compare our method with the following benchmarks:

1-NN-s and SVM-s : Basic 1-NN and linear SVM classifiers trained only on the source domain.

GFK-PLS (Gong et al., 2012) : The geodesic flow kernel algorithm where partial least squares (PLS) implementation is used to initialize the source subspace. Results are evaluated on kernel-NNs.

SA (Fernando et al., 2013) : This is the subspace alignment algorithm. Results are evaluated using 1-NN.

CORAL (Sun et al., 2016) : The correlation alignment algorithm that uses a linear SVM on the similarity matrix formed by correlation matching.

5.3.1 Office+Caltech10 (Unsupervised)

We follow the original protocol provided by Gong et al. (2012) on Office+Caltech10 dataset. Note that several baselines, determine the best dimensionality per domain to achieve their maximum accuracies on SURF features. We observed that a dimensionality in the range [80,120] provides consistent results for our solution using SURF features. For VGG features we empirically found the dimensionality of 20 suits best for the compared DA-SL algorithms.

Table 6: Unsupervised domain adaptation results using SURF features on Office+Caltech10 (Gong et al., 2012) dataset with the evaluation setup of Gong et al. (2012).The best (in bold blue), the second best (in blue).

	A→W	A→D	A→C	W→A	W→D	W→C	D→A	D→W	D→C	C→A	C→W	C→D	**Avg.**
1-NN-s	23.1	22.3	20.0	13.8	40.6	12.2	23.0	51.7	19.9	21.0	19.0	23.6	24.2
SVM-s	25.6	33.4	35.9	32.1	78.9	25.2	34.6	70.2	31.2	43.8	30.5	40.3	40.1
GFK-PLS	35.7	35.1	37.9	35.5	71.2	29.3	36.2	79.1	32.7	40.4	35.8	41.1	42.5
SA	38.6	37.6	35.3	37.4	80.3	32.3	38.0	83.6	32.4	39.0	36.8	39.6	44.2
CORAL	38.7	38.3	**40.3**	37.8	**84.9**	**34.6**	38.1	**85.9**	34.2	47.2	39.2	40.7	**46.7**
ILS (1-NN)	**40.6**	**41.0**	37.1	**39.0**	78.7	34.2	**38.9**	79.1	**36.9**	**48.6**	**42.0**	**44.1**	**46.7**

Table 7: Unsupervised domain adaptation results using VGG-FC6 features on Office+Caltech10 (Gong et al., 2012) dataset with the evaluation setup of Gong et al. (2012).The best (in bold blue), the second best (in blue).

	A→W	A→D	A→C	W→A	W→D	W→C	D→A	D→W	D→C	C→A	C→W	C→D	**Avg.**
1-NN-s	60.9	52.3	70.1	66.4	91.3	60.2	57.0	86.7	48.0	81.9	65.9	55.6	66.4
SVM-s	63.1	51.7	74.2	73.3	94.2	68.2	58.7	91.8	55.5	86.7	74.8	61.5	71.1
GFK-PLS	74.1	63.5	77.7	81.1	**96.6**	73.5	69.9	92.4	64.0	86.2	76.5	66.5	76.8
SA	76.0	64.9	77.1	80.2	94.2	71.9	69.0	90.5	62.3	83.9	76.0	66.2	76.0
CORAL	74.8	67.1	**79.0**	82.3	96.0	75.9	75.8	**94.6**	64.7	**89.4**	77.6	67.6	78.7
ILS (1-NN)	**82.4**	**72.5**	78.9	**87.2**	89.3	**79.9**	**79.2**	94.2	**66.5**	87.6	**84.4**	**73.0**	**81.3**

Table 8: Unsupervised domain adaptation results using VGG-FC7 features on Office+Caltech10 (Gong et al., 2012) dataset with the evaluation setup of Gong et al. (2012).The best (in bold blue), the second best (in blue).

	A→W	A→D	A→C	W→A	W→D	W→C	D→A	D→W	D→C	C→A	C→W	C→D	**Avg.**
1-NN-s	64.0	50.8	72.6	67.8	88.8	64.2	61.2	88.2	52.8	82.6	65.3	54.9	67.8
SVM-s	68.0	51.8	76.2	74.6	93.0	70.6	58.7	91.2	56.0	86.7	74.8	61.3	71.9
GFK-PLS	74.0	57.6	76.6	76.0	92.9	69.5	67.5	**91.9**	62.9	84.1	73.6	63.4	74.2
SA	75.0	60.7	76.2	76.4	94.0	69.0	66.0	89.5	59.4	82.6	73.6	63.2	73.8
CORAL	71.8	61.3	**78.6**	82.0	**94.6**	73.7	71.2	93.5	63.0	**88.6**	76.0	63.8	76.5
ILS (1-NN)	**80.9**	**71.3**	78.4	**86.7**	88.2	**76.3**	**76.5**	91.8	**66.2**	87.1	**80.1**	**67.1**	**79.2**

Table. 6, Table. 7 and Table. 8 present the unsupervised setting results using the SURF, VGG-FC6 and VGG-FC7 features. For all feature types, our solution yields the best overall performance as well as the best performance in 8 domain transformations out of 12. Similarly to the semi-supervised experiments we notice the VGG-FC7 is less favorable for DA algorithms. Nevertheless, our solution achieves the highest overall classification accuracies for both FC6 and FC7 features.

Learned Transformations as Subspace Representations: We consider both GFK (Gong et al., 2012) and SA (Fernando et al., 2013) as DA-SL algorithms. Both these methods make use of PCA subspaces to adapt the domains. To the best of our knowledge, there exists no through studies claiming that PCA is the method of choice to for GFK and SA. As a matter of fact, Gong *et al*. show that the performance of GFK can be improved if PLS[9] algorithm is employed to define the source subspace (Gong et al., 2012). Whether PCA or PLS is used to define the subspaces, identification of the subspaces is disjoint from the domain adaptation technique in GFK and SA. In contrast, transformations in the ILS algorithm are linked to the adaptation process. This makes a curious

9. Despite using labeled data, this method falls under the unsupervised setting since it does not use the labeled target data.

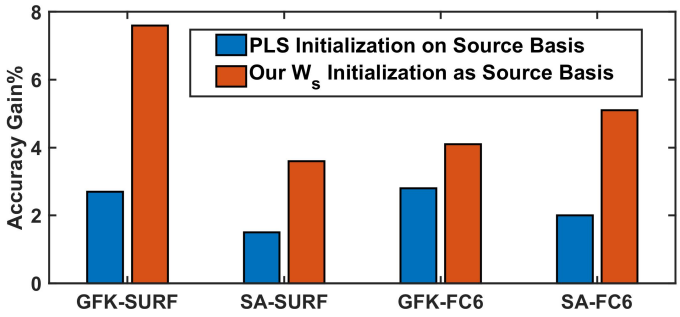

Figure 5: The accuracy gain on Office-Caltech dataset for GFK (Gong et al., 2012) and SA (Fernando et al., 2013) when their initial PCA subspaces are replaced with PLS and our W_s transformation matrices.

mind wondering whether the learned transformations in the ILS algorithm capture better structures for adaptation or not. We empirically show that this is indeed the case by using the learned $\boldsymbol{W}_s$ as the source subspace in GFK and SA.

Figure 5 compares the accuracy gains over PCA spaces by using PLS and our $\boldsymbol{W}_s$ initialization. It is clear that the highest classification accuracy gain is obtained by our $\boldsymbol{W}_s$ initialization. This proves that $\boldsymbol{W}_s$ is capable to learn a more favorable subspace representation for domain adaptation.

5.3.2 PIE-Multiview Faces

The PIE Multiview dataset includes face images of 67 individuals captured from different views, illumination conditions, and expressions. In this experiment, we use the views $C27$ (looking forward) as the source domain and $C09$ (looking down), and the views $C05$, $C37$, $C02$, $C25$ (looking towards left in an increasing angle, see Fig. 6) as target domains. We expect the face inclination angle to reflect the complexity of transfer learning. We normalize the images to $32\times$ 32 pixels and use the vectorized gray-scale images as features. Empirically, we observed that the GFK (Gong et al., 2012) and SA (Fernando et al., 2013) reach better performances if the features are normalized to have unit ℓ_2 norm. We therefore use ℓ_2 normalized features in our evaluations. The dimensionality of the subspaces for all the subspace based methods (*i.e.*, Gong et al. (2012); Fernando et al. (2013)) including ours is 100.

Table. 9 lists the classification accuracies with increasing angle of inclination. Our solution attains the best average accuracy as well as the best scores for 4 views and the second best for the $C09$. With the increasing camera angle, the feature structure changes up to a certain extent. In other words, the features become heterogeneous. However, our algorithm boosts the accuracies even under such challenging conditions.

6. Parameter Sensitivity and Orthogonality

In all the above experiments, we keep $\lambda = 1$ (see Eq. 1). To analyze the sensitivity of our method to the changes in parameter λ, we performed an experiment using the unsupervised protocol. This is because the statistical loss plays a significant role in establishing the correspondence between

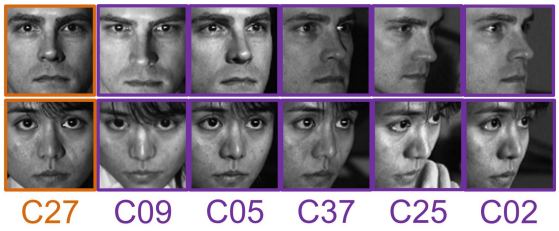

Figure 6: Two instances of the PIE-Multiview face data. Here, the view from $C27$ is used as the source domain. Remaining views are considered to be the target for each transformation.

Table 9: PIE-Multiview results. The variation of performance w.r.t. face orientations when frontal face images are considered as the source domain.

camera pose→	C09	C05	C37	C25	C02	**Avg.**
1-NN-s	92.5	55.7	28.5	14.8	11.0	40.5
SVM-s	87.8	65.0	35.8	15.7	16.7	44.2
GFK-PLS	92.5	74.0	32.1	14.1	12.3	45.0
SA	**97.9**	85.9	47.9	16.6	13.9	52.4
CORAL	91.4	74.8	35.3	13.4	13.2	45.6
ILS (1-NN)	96.6	**88.3**	**72.9**	**28.4**	**34.8**	**64.2**

the source and the target in the unsupervised DA. We consider two random splits from each of the Office+Caltech10 dataset along VGG-FC6 features here.

Our results are shown in Fig. 7. When $\lambda = 0$, no statistical loss term is considered. It is clear that for this case the performance drops considerably. For other values of λ, the performance is superior and there is little variation in performance. In other words, our method remains robust.

We further investigate the benefit of orthogonality constraint on $\boldsymbol{W}_s$ and $\boldsymbol{W}_t$ against free-form and unconstrained transformations. Using the orthogonality constraint provides a considerable performance gain as shown in Fig. 7. While orthogonality makes the optimization more complicated, it seems it guides the learning to better uncovering the form of adaptation.

Conclusion

In this paper, we proposed a solution for both semi-supervised and unsupervised Domain Adaptation (DA) problems. Our solution learns a latent space in which domain discrepancies are minimized. We showed that such a latent space can be obtained by **1.** minimizing a notion of discriminatory power over the available labeled data while simultaneously **2.** matching statistical properties across the domains. To determine the latent space, we modeled the learning problem as a minimization problem on Riemannian manifolds and solved it using optimization techniques on matrix manifolds.

Empirically, we showed that the proposed method outperformed state-of-the-art DA solutions in semi-supervised and unsupervised settings using both hand-crafted and deep-net features. As a future work, we plan to tackle two related problems, namely multiple source DA and domain generalization (Ghifary et al., 2016).

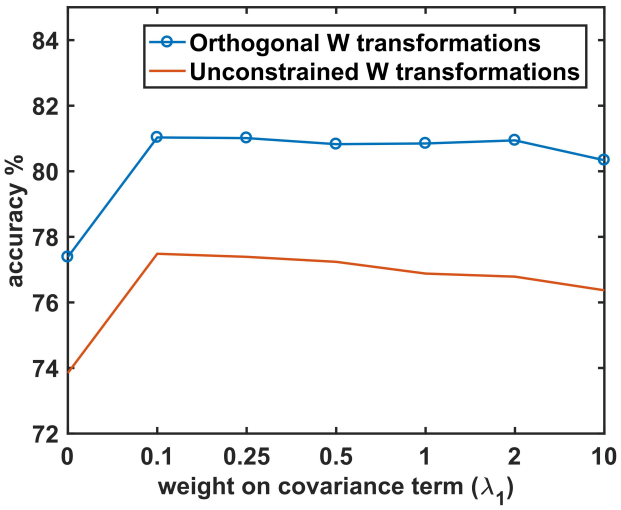

Figure 7: Sensitivity to λ in the unsupervised DA.

References

P-A Absil, Robert Mahony, and Rodolphe Sepulchre. *Optimization algorithms on matrix manifolds*. Princeton University Press, 2009.

Mahsa Baktashmotlagh, Mehrtash Harandi, and Mathieu Salzmann. Distribution-matching embedding for visual domain adaptation. *Journal of Machine Learning Research*, 17(108):1–30, 2016.

Herbert Bay, Tinne Tuytelaars, and Luc Van Gool. Surf: Speeded up robust features. In *European conference on computer vision*, pages 404–417. Springer, 2006.

Karsten Borgwardt, Arthur Gretton, Malte J. Rasch, Hans-Peter Kriegel, Bernhard Schoelkopf, and Alexander Smola. Integrating structured biological data by kernel maximum mean discrepancy. *Bioinformatics*, 22:49–57, 2006.

N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre. Manopt, a Matlab toolbox for optimization on manifolds. *Journal of Machine Learning Research*, 15:1455–1459, 2014. URL http://www.manopt.org.

Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.

Mike Brookes. The matrix reference manual. *Imperial College London*, 2005.

Qiang Chen, Junshi Huang, Rogerio Feris, Lisa M Brown, Jian Dong, and Shuicheng Yan. Deep domain adaptation for describing people based on fine-grained clothing attributes. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5315–5324, 2015.

A. Cherian, S. Sra, A. Banerjee, and N. Papanikolopoulos. Jensen-bregman logdet divergence with application to efficient similarity search for covariance matrices. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(9):2161–2174, Sept 2013.

Hal Daumé III, Abhishek Kumar, and Avishek Saha. Frustratingly easy semi-supervised domain adaptation. In *Proceedings of the 2010 Workshop on Domain Adaptation for Natural Language Processing*, pages 53–59, 2010.

Jason V. Davis, Brian Kulis, Prateek Jain, Suvrit Sra, and Inderjit S. Dhillon. Information-theoretic metric learning. In *ICML*, pages 209–216, Corvalis, Oregon, USA, June 2007.

Lixin Duan, Dong Xu, and Ivor W. Tsang. Learning with augmented features for heterogeneous domain adaptation. In *Proc. Int. Conference on Machine Learning (ICML)*, pages 711–718, June 2012.

Alan Edelman, Tomás A Arias, and Steven T Smith. The geometry of algorithms with orthogonality constraints. *SIAM journal on Matrix Analysis and Applications*, 20(2):303–353, 1998.

B. Fernando, A. Habrard, M. Sebban, and T. Tuytelaars. Unsupervised visual domain adaptation using subspace alignment. In *Proc. Int. Conference on Computer Vision (ICCV)*, pages 2960–2967, 2013.

Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *Proc. Int. Conference on Machine Learning (ICML)*, pages 1180–1189, 2015.

M. Ghifary, D. Balduzzi, W. B. Kleijn, and M. Zhang. Scatter component analysis: A unified framework for domain adaptation and domain generalization. *IEEE Trans. Pattern Analysis and Machine Intelligence*, PP(99):1–1, 2016.

B. Gong, Y. Shi, F. Sha, and K. Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2066–2073, 2012.

R. Gopalan, Ruonan Li, and R. Chellappa. Domain adaptation for object recognition: An unsupervised approach. In *Proc. Int. Conference on Computer Vision (ICCV)*, pages 999–1006, 2011.

Victor Guillemin and Alan Pollack. *Differential topology*, volume 370. American Mathematical Soc., 2010.

Mehrtash Tafazzoli Harandi, Mathieu Salzmann, and Richard I. Hartley. Dimensionality reduction on SPD manifolds: The emergence of geometry-aware methods. *CoRR*, abs/1605.06182, 2016.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, June 2016.

Judy Hoffman, Erik Rodner, Jeff Donahue, Kate Saenko, and Trevor Darrell. Efficient learning of domain-invariant image representations. In *International Conference on Learning Representations*, 2013.

Judy Hoffman, Erik Rodner, Jeff Donahue, Brian Kulis, and Kate Saenko. Asymmetric and category invariant feature transformations for domain adaptation. *Int. Journal of Computer Vision*, 109(1): 28–41, 2014.

Yao-Hung Hubert Tsai, Yi-Ren Yeh, and Yu-Chiang Frank Wang. Learning cross-domain landmarks for heterogeneous domain adaptation. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5081–5090, June 2016.

Ioan Mackenzie James. *The topology of Stiefel manifolds*, volume 24. Cambridge University Press, 1976.

Martin Köstinger, Martin Hirzer, Paul Wohlhart, Peter M Roth, and Horst Bischof. Large scale metric learning from equivalence constraints. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 2288–2295, 2012.

Anjali Krishnan, Lynne J Williams, Anthony Randal McIntosh, and Hervé Abdi. Partial least squares (pls) methods for neuroimaging: a tutorial and review. *Neuroimage*, 56(2):455–475, 2011.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, pages 1097–1105, 2012.

B. Kulis, K. Saenko, and T. Darrell. What you saw is not what you get: Domain adaptation using asymmetric kernel transforms. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1785–1792, June 2011.

John M Lee. Smooth manifolds. In *Introduction to Smooth Manifolds*, pages 1–29. Springer, 2003.

M. Long, J. Wang, Y. Cao, J. Sun, and P. S. Yu. Deep learning of transferable representation for scalable domain adaptation. *IEEE Transactions on Knowledge and Data Engineering*, 28(8): 2027–2040, Aug 2016.

S. J. Pan, I. W. Tsang, J. T. Kwok, and Q. Yang. Domain adaptation via transfer component analysis. *IEEE Transactions on Neural Networks*, 22(2):199–210, 2011.

V. M. Patel, R. Gopalan, R. Li, and R. Chellappa. Visual domain adaptation: A survey of recent advances. *IEEE Signal Processing Magazine*, 32(3):53–69, 2015.

Xavier Pennec, Pierre Fillard, and Nicholas Ayache. A riemannian framework for tensor computing. *Int. Journal of Computer Vision*, 66(1):41–66, 2006.

Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *Int. Journal of Computer Vision*, 115 (3):211–252, 2015.

Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *Proc. European Conference on Computer Vision (ECCV)*, pages 213–226, 2010.

Xiaoxiao Shi, Qi Liu, Wei Fan, S Yu Philip, and Ruixin Zhu. Transfer learning on heterogenous feature spaces via spectral transformation. In *2010 IEEE international conference on data mining*, pages 1049–1054, 2010.

Hidetoshi Shimodaira. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of Statistical Planning and Inference*, 90(2):227 – 244, 2000.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.

A. Torralba and A. A. Efros. Unbiased look at dataset bias. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1521–1528, 2011.

Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014.

Chang Wang and Sridhar Mahadevan. Heterogeneous domain adaptation using manifold alignment. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume Two*, IJCAI'11, pages 1541–1546, 2011.

Appendix

In this appendix, we provide more details on the product geometry of the problem discussed in § 2 and also the form of gradients required to perform Riemannian optimization.

A. Product Topology

As the constraints of the optimization problem depicted in Eq. 1 are indeed Riemannian manifolds, the whole set of constraints can be given a Riemannian structure through the concept of product topology. In particular, the constraints can be modeled as

$$\mathcal{M}_{prod.} = \mathrm{St}(p,s) \times \mathrm{St}(p,t) \times \mathcal{S}^{p}_{++} \times \mathbb{R}^{N_p}, \tag{13}$$

The tangent space of such a product topology (Lee, 2003; Guillemin and Pollack, 2010) could be written as,

$$\mathcal{T}_{(\boldsymbol{W}_s,\boldsymbol{W}_t,\boldsymbol{M},\boldsymbol{\epsilon})}\mathcal{M}_{prod.} = T_{\boldsymbol{W}_s}\mathrm{St}(p,s) \times T_{\boldsymbol{W}_t}\mathrm{St}(p,t) \times T_{\boldsymbol{M}}\mathcal{S}^{p}_{++} \times T_{\boldsymbol{\epsilon}}\mathbb{R}^{N_p}. \tag{14}$$

In Table 10, the metric and, the form of Riemannian gradient and the retraction for $\mathcal{M}_{prod.}$ are provided.

B. Derivations

We recall that the cost function depicted in Eq. 1 consists of two parts, namely $\mathcal{L}_d$ and $\mathcal{L}_u$. Here, $\mathcal{L}_d$ is a measure of dissimilarity between labeled samples. The term $\mathcal{L}_u$ quantifies a notion of statistical difference between the source and target samples in the latent space. We provide the gradients of Eq. 1 with respect to its parameters below. This, as discussed in § 3 is required for Riemannian optimization.

Derivative of soft-margin ℓ_β

We recall that $\mathcal{L}_d$ has the following form,

$$\mathcal{L}_d = \frac{1}{N_p}\sum_{k=1}^{N_p} \ell_\beta\big(\boldsymbol{M}, y_k, \boldsymbol{z}_{1,k} - \boldsymbol{z}_{2,k}, 1 + y_k\epsilon_k\big) + r(\boldsymbol{M}) + \frac{1}{N_p}\sqrt{\sum \epsilon_k^2}\,, \tag{15}$$

with

$$\ell_\beta\big(\boldsymbol{M}, y, \boldsymbol{x}, u\big) = \frac{1}{\beta}\log\Big(1 + \exp\big(\beta y(\boldsymbol{x}^T\boldsymbol{M}\boldsymbol{x} - u)\big)\Big). \tag{16}$$

In Eq. 15, y_k denotes whether the k-th pair is similar or dissimilar (*i.e.*, $y_k = +1$ if $\boldsymbol{z}_{1,k}$ and $\boldsymbol{z}_{2,k}$ are from the same class and $y_k = -1$ otherwise).

For the sake of discussion, assume $\boldsymbol{z}_{1,k}$ and $\boldsymbol{z}_{2,k}$ are embedded from the source and target domains, respectively. That is $\boldsymbol{z}_{1,k} = \boldsymbol{W}_s^T\boldsymbol{x}_i^s$ and $\boldsymbol{z}_{2,k} = \boldsymbol{W}_t^T\boldsymbol{x}_j^t$. By expanding ℓ_β for such a pair, we get

Table 10: Riemannian metric, gradient and retraction on the proposed Product Manifold in Eq. 13. The Riemannian metrics g_{ν_s} and g_{ν_t} are respectively defined on the Stiefel manifolds of $\boldsymbol{W}_s$ and $\boldsymbol{W}_t$. Furthermore, the Riemannian metrics g_S and g_E are respectively on the SPD manifold and the Euclidean manifold. As we have used in the main text, ξ and ς are elements from the tangent spaces of the corresponding manifolds. Here, $\mathrm{uf}(\boldsymbol{A}) = \boldsymbol{A}(\boldsymbol{A}^T\boldsymbol{A})^{-1/2}$, which yields an orthogonal matrix, $\mathrm{sym}(\boldsymbol{A}) = \frac{1}{2}(\boldsymbol{A} + \boldsymbol{A}^T)$ and $\mathrm{expm}(\cdot)$ denotes the matrix exponential.

	$\mathcal{M}_{prod.}$
Matrix representation	$(\boldsymbol{W}_s, \boldsymbol{W}_t, \boldsymbol{M}, \boldsymbol{\epsilon})$
Riemannian metric	$g_{\nu_s}(\varsigma_s, \xi_s) + g_{\nu_t}(\varsigma_t, \xi_t) + g_S(\varsigma_M, \xi_M) + g_E(\varsigma_E, \xi_E)$
Riemannian gradient	$\Big(\nabla_{\boldsymbol{W}_s}(f) - \boldsymbol{W}_s\mathrm{sym}\Big(\boldsymbol{W}_s^T\nabla_{\boldsymbol{W}_s}(f)\Big), \nabla_{\boldsymbol{W}_t}(f) - \boldsymbol{W}_t\mathrm{sym}\Big(\boldsymbol{W}_t^T\nabla_{\boldsymbol{W}_t}(f)\Big), \boldsymbol{M}\mathrm{sym}\Big(\nabla_{\boldsymbol{M}}(f)\Big)\boldsymbol{M}, \nabla_{\boldsymbol{\epsilon}}(f)\Big)$
Retraction	$\big(\mathrm{uf}(\boldsymbol{W}_s + \xi_s), \mathrm{uf}(\boldsymbol{W}_t + \xi_t), \boldsymbol{M}^{\frac{1}{2}}\,\mathrm{expm}(\boldsymbol{M}^{-\frac{1}{2}}\xi_M\boldsymbol{M}^{-\frac{1}{2}})\boldsymbol{M}^{\frac{1}{2}}, \mathbf{I}_p\big)$

$$\ell_\beta(\boldsymbol{M}, y_k, \boldsymbol{z}_{1,k} - \boldsymbol{z}_{2,k}, 1 + y_k\epsilon_k) = \frac{1}{\beta}\log(1 + \exp(\beta y_k((\boldsymbol{z}_{1,k} - \boldsymbol{z}_{2,k})^T\boldsymbol{M}(\boldsymbol{z}_{1,k} - \boldsymbol{z}_{2,k}) - 1 - y_k\epsilon_k))) \tag{17}$$

To simplify the presentation, we define $d(\boldsymbol{M}, \boldsymbol{W}_s, \boldsymbol{W_t}) = (\boldsymbol{z}_{1,k} - \boldsymbol{z}_{2,k})^T\boldsymbol{M}(\boldsymbol{z}_{1,k} - \boldsymbol{z}_{2,k})$ and $r = \exp(\beta y_k((\boldsymbol{z}_{1,k} - \boldsymbol{z}_{2,k})^T\boldsymbol{M}(\boldsymbol{z}_{1,k} - \boldsymbol{z}_{2,k}) - y_k\epsilon_k - 1))$. We provide the gradients of Eq. 17 with respect to $\boldsymbol{M}$, $\boldsymbol{W_t}$, $\boldsymbol{W_s}$ and the slack ϵ_k below.

DERIVATIVE W.R.T. $\boldsymbol{M}$

$$\begin{aligned}
\nabla_{\boldsymbol{M}}\ell_\beta &= \frac{y_k r}{(1+r)}\nabla_{\boldsymbol{M}}d(\boldsymbol{M}) \\
&= \frac{y_k r}{(1+r)}(\boldsymbol{W}_s^T\boldsymbol{x}_i^s - \boldsymbol{W}_t^T\boldsymbol{x}_j^t)(\boldsymbol{x}_i^{sT}\boldsymbol{W}_s - \boldsymbol{x}_j^{tT}\boldsymbol{W}_t) \\
&= y_k(1 + r^{-1})^{-1}(\boldsymbol{W}_s^T\boldsymbol{x}_i^s - \boldsymbol{W}_t^T\boldsymbol{x}_j^t)(\boldsymbol{x}_i^{sT}\boldsymbol{W}_s - \boldsymbol{x}_j^{tT}\boldsymbol{W}_t).
\end{aligned} \tag{18}$$

DERIVATIVE W.R.T. $\boldsymbol{W}_s$ (OR W.R.T. $\boldsymbol{W}_t$)

$$\begin{aligned}
\nabla_{\boldsymbol{W}_s}\ell_\beta &= \frac{y_k r}{(1+r)}\nabla_{\boldsymbol{W}_s}d(\boldsymbol{W}_s) && (19)\\
&= 2\frac{y_k r}{(1+r)}\boldsymbol{x}_i^s(\boldsymbol{x}_i^{sT}\boldsymbol{W}_s - \boldsymbol{x}_j^{tT}\boldsymbol{W}_t)\boldsymbol{M} \\
&= 2y_k(1 + r^{-1})^{-1}\boldsymbol{x}_i^s(\boldsymbol{x}_i^{sT}\boldsymbol{W}_s - \boldsymbol{x}_j^{tT}\boldsymbol{W}_t)\boldsymbol{M}. && (20)
\end{aligned}$$

For the case where both the pair instances are from the same domain (*i.e.* $\boldsymbol{z}_{1,k} = \boldsymbol{W}_s^T\boldsymbol{x}_i^s$ and $\boldsymbol{z}_{2,k} = \boldsymbol{W}_s^T\boldsymbol{x}_j^s$), it could be shown that,

$$\nabla_{\boldsymbol{W}_s}d(\boldsymbol{W}_s) = 2y_k(\boldsymbol{x}_i^s - \boldsymbol{x}_j^s)(\boldsymbol{x}_i^{sT} - \boldsymbol{x}_j^{sT})\boldsymbol{W}_s\boldsymbol{M}. \tag{21}$$

Considering Eq. 19 and Eq. 21,

Table 11: Summarized Gradients. Note we use $r = \exp(\beta y_k(d(\boldsymbol{M}) - y_k\epsilon_k - 1))$.

$\nabla_{\boldsymbol{W}_s}\ell_\beta\ ; x_i^s \in \mathbb{R}^s, x_j^t \in \mathbb{R}^t$	$2y_k(1+r^{-1})^{-1}\boldsymbol{x}_i^s(\boldsymbol{x}_i^{sT}\boldsymbol{W}_s - \boldsymbol{x}_j^{t\,T}\boldsymbol{W}_t)\boldsymbol{M}$
$\nabla_{\boldsymbol{W}_s}\ell_\beta\ ; x_i^s \in \mathbb{R}^s, x_j^s \in \mathbb{R}^s$	$2y_k(1+r^{-1})^{-1}(\boldsymbol{x}_i^s - \boldsymbol{x}_j^s)(\boldsymbol{x}_i^{sT} - \boldsymbol{x}_j^{sT})\boldsymbol{W}_s\boldsymbol{M}$
$\nabla_{\boldsymbol{W}_t}\ell_\beta\ ; x_i^s \in \mathbb{R}^s, x_j^t \in \mathbb{R}^t$	$2y_k(1+r^{-1})^{-1}\boldsymbol{x}_j^t(\boldsymbol{x}_j^{t\,T}\boldsymbol{W}_t - \boldsymbol{x}_i^{sT}\boldsymbol{W}_s)\boldsymbol{M}$
$\nabla_{\boldsymbol{W}_t}\ell_\beta\ ; x_i^t \in \mathbb{R}^t, x_j^t \in \mathbb{R}^t$	$2y_k(1+r^{-1})^{-1}(\boldsymbol{x}_i^t - \boldsymbol{x}_j^t)(\boldsymbol{x}_i^{t\,T} - \boldsymbol{x}_j^{t\,T})\boldsymbol{W}_t\boldsymbol{M}$
$\nabla_{\boldsymbol{M}}\ell_\beta$	$2y_k(1+r^{-1})^{-1}(\boldsymbol{W}_s^T\boldsymbol{x}_i^s - \boldsymbol{W}_t^T\boldsymbol{x}_j^t)(\boldsymbol{x}_i^{sT}\boldsymbol{W}_s - \boldsymbol{x}_j^{t\,T}\boldsymbol{W}_t)$
$\nabla_{v_k}\ell_\beta$	$-e^{v_k}(1+r^{-1})^{-1}$
$\nabla_{\boldsymbol{W}_s}\mathcal{L}_u$	$\frac{1}{p}\boldsymbol{\Sigma}_s\boldsymbol{W}_s\Big(2\big(\boldsymbol{W}_s^T\boldsymbol{\Sigma}_s\boldsymbol{W}_s + \boldsymbol{W}_t^T\boldsymbol{\Sigma}_t\boldsymbol{W}_t\big)^{-1} - \big(\boldsymbol{W}_s^T\boldsymbol{\Sigma}_s\boldsymbol{W}_s\big)^{-1}\Big)$

$$\begin{aligned}\nabla_{\boldsymbol{W}_s}\ell_\beta &= 2\frac{y_k r}{(1+r)}(\boldsymbol{x}_i^s - \boldsymbol{x}_j^s)(\boldsymbol{x}_i^{sT} - \boldsymbol{x}_j^{sT})\boldsymbol{W}_s\boldsymbol{M}.\\ &= 2y_k(1+r^{-1})^{-1}(\boldsymbol{x}_i^s - \boldsymbol{x}_j^s)(\boldsymbol{x}_i^{sT} - \boldsymbol{x}_j^{sT})\boldsymbol{W}_s\boldsymbol{M}. \end{aligned} \tag{22}$$

DERIVATIVE W.R.T. A SLACK VARIABLE ϵ_k.

The slacks by origin are non-negative. To avoid using a non-negative constraint we make the substitution $\epsilon_k = e^{v_k}$ to Eq. 17.

$$\therefore \ell_\beta = \frac{1}{\beta}\log(1 + \exp(\beta y_k(d(\boldsymbol{M}, \boldsymbol{W}_s, \boldsymbol{W_t}, v_k) - 1 - y_k e^{v_k}))) \tag{23}$$

The derivative of Eq. 23 w.r.t. v_k,

$$\begin{aligned}\nabla_{v_k}\ell_\beta &= \frac{-e^{v_k}\exp(\beta y_k(d(\boldsymbol{M}, \boldsymbol{W}_s, \boldsymbol{W_t}, v_k) - 1 - y_k e^{v_k}))}{(1 + \exp(\beta y_k(d(\boldsymbol{M}, \boldsymbol{W}_s, \boldsymbol{W_t}, v_k) - 1 - y_k e^{v_k}))}\\ &= \frac{-e^{v_k} r}{(1+r)} = -e^{v_k}(1+r^{-1})^{-1}\end{aligned} \tag{24}$$

Derivative of Statistical loss $\mathcal{L}_u$

The statistical loss (*i.e.* unsupervised loss) in Eq. 1 is defined using the stein divergence δ_s. We have,

$$\begin{aligned}\mathcal{L}_u &= \frac{1}{p}\delta_s\big(\boldsymbol{W}_s^T\boldsymbol{\Sigma}_s\boldsymbol{W}_s, \boldsymbol{W}_t^T\boldsymbol{\Sigma}_t\boldsymbol{W}_t\big)\\ &= \frac{1}{p}\left\{\log\det\left(\frac{\boldsymbol{W}_s^T\boldsymbol{\Sigma}_s\boldsymbol{W}_s + \boldsymbol{W}_t^T\boldsymbol{\Sigma}_t\boldsymbol{W}_t}{2}\right) - \frac{1}{2}\log\det(\boldsymbol{W}_s^T\boldsymbol{\Sigma}_s\boldsymbol{W}_s\boldsymbol{W}_t^T\boldsymbol{\Sigma}_t\boldsymbol{W}_t)\right\}\\ &= \frac{1}{p}\left\{\log\det\left(\frac{\boldsymbol{W}_s^T\boldsymbol{\Sigma}_s\boldsymbol{W}_s + \boldsymbol{W}_t^T\boldsymbol{\Sigma}_t\boldsymbol{W}_t}{2}\right) - \frac{1}{2}\log\det(\boldsymbol{W}_s^T\boldsymbol{\Sigma}_s\boldsymbol{W}_s) - \frac{1}{2}\log\det(\boldsymbol{W}_t^T\boldsymbol{\Sigma}_t\boldsymbol{W}_t)\right\},\end{aligned} \tag{25}$$

where $\boldsymbol{\Sigma}_s$ and $\boldsymbol{\Sigma}_t$ are respectively the source and target domain covariances. Making use of 2.11 of Brookes (2005), we yield

$$\nabla_{\boldsymbol{W}_s} \log\det(\boldsymbol{W}_s^T \boldsymbol{\Sigma}_s \boldsymbol{W}_s) = 2\boldsymbol{\Sigma}_s \boldsymbol{W}_s (\boldsymbol{W}_s^T \boldsymbol{\Sigma}_s \boldsymbol{W}_s)^{-1}. \tag{26}$$

The derivative of Eq. 25 with respect to $\boldsymbol{W}_s$ (or similarly for $\boldsymbol{W}_t$[10]) can be obtained,

$$\begin{aligned}
\nabla_{\boldsymbol{W}_s} \mathcal{L}_u &= \frac{1}{p} \boldsymbol{\Sigma}_s \boldsymbol{W}_s \left\{ \left(\frac{\boldsymbol{W}_s^T \boldsymbol{\Sigma}_s \boldsymbol{W}_s + \boldsymbol{W}_t^T \boldsymbol{\Sigma}_t \boldsymbol{W}_t}{2} \right)^{-1} - (\boldsymbol{W}_s^T \boldsymbol{\Sigma}_s \boldsymbol{W}_s)^{-1} \right\} \\
&= \frac{1}{p} \boldsymbol{\Sigma}_s \boldsymbol{W}_s \left\{ 2\Big(\boldsymbol{W}_s^T \boldsymbol{\Sigma}_s \boldsymbol{W}_s + \boldsymbol{W}_t^T \boldsymbol{\Sigma}_t \boldsymbol{W}_t \Big)^{-1} - (\boldsymbol{W}_s^T \boldsymbol{\Sigma}_s \boldsymbol{W}_s)^{-1} \right\}
\end{aligned} \tag{27}$$

The derivatives are summarized in Table 11.

10. The Stein divergence is symmetric over its two arguments.